

CONF-790510--1

LA-UR-79-897

TITLE: IMAGE RESTORATION WITH A LOCALLY VARIABLE
WIENER FILTER

AUTHOR(S): Martin J. Lahart

SUBMITTED TO: SPIE Huntsville Electro-Optical Technical
Symposium and Workshop, May 22-25, 1979,
Huntsville, AL.

NOTICE

This report was prepared as an account of work sponsored by the United States Government. Neither the United States nor the United States Department of Energy, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness or usefulness of any information, apparatus, product or process disclosed, or represents that its use would not infringe privately owned rights.

By acceptance of this article for publication, the publisher recognizes the Government's (license) rights in any copyright and the Government and its authorized representatives have unrestricted right to reproduce in whole or in part said article under any copyright secured by the publisher.

The Los Alamos Scientific Laboratory requests that the publisher identify this article as work performed under the auspices of the USERDA.


los alamos
scientific laboratory
of the University of California
LOS ALAMOS, NEW MEXICO 87545

An Affirmative Action/Equal Opportunity Employer

MASTER

Form No. 846
St. No. 2829
1/75

UNITED STATES
ENERGY RESEARCH AND
DEVELOPMENT ADMINISTRATION
CONTRACT W-7405-ENG. 36

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

FMB

IMAGE RESTORATION WITH A LOCALLY VARIABLE WIENER FILTER*

M. J. Lahart
 University of California
 Los Alamos Scientific Laboratory
 Los Alamos, New Mexico 87545

Abstract

The parametric Wiener filter is often used to deblur images that are relatively noise-free. If noise is more severe, the restored image may be obscured by a granular pattern that results when the noise is subjected to the deblurring filter. This effect may be reduced by using a larger noise parameter, but this leads to a restoration that is less sharp. We describe how the noise parameter may be varied from pixel to pixel, so that it is larger only where noise is greater. Pixels with low signal-to-noise ratios are identified by a thresholding process and by comparison with nearest neighbors. The effects of the estimated Wiener spectra on the restored image are discussed.

Introduction

The application of small computers to image processing has made the development of local processing techniques important. When memory is limited, techniques are preferred that process small portions of images at a time so that an entire image might be processed in a serial manner. However, one of the most important tools in image restoration, that of spatial filtering, is generally global in nature--Fourier transforms of entire images must be taken before filtering can be performed by multiplication in the frequency domain. If images are divided into smaller subimages in an attempt to approach serial processing, the influence of the edges of the segments becomes important. High-pass filtered images exhibit ripples extending away from the edges if high-frequency boosting is mild; if it is severe, the entire filtered subimage may be obliterated by artifacts. When larger images are processed, artifacts are less severe, but time consuming writes to disk are necessary to interchange rows and columns.

Computations are performed in the Fourier domain, because the convolution matrix is diagonal there. In the direct domain, the solution of the equation that describes convolution is often thought to require the inversion of a matrix with an extremely large number of elements, sometimes as large as the square of the number of pixels in the image. We will show that images can be deconvolved in the direct domain with the inversion of a small enough matrix so that computations can be performed easily. The value of each pixel in the object is estimated separately, and only a few pixels in its neighborhood are used in the estimation. Results appear the same as global restorations produced in the Fourier domain.

Besides facilitating computations, local filtering offers a flexibility that is not attainable with global techniques. Filter parameters can be varied from pixel to pixel so that spatially variant point spread functions or signal-to-noise ratios can be accommodated in restoration algorithms. We will show by example how parameters can be varied to adjust for a spatially variant signal-to-noise ratio. Our example uses a binary filter in which one of only two sets of filter parameters are used in the restoration of each pixel. The technique can be extended to permit choices from much larger sets of parameters.

Examination of Global Techniques

We will compare local and global restoration techniques by blurring a particular image, adding Gaussian random noise to it when it is appropriate to add noise, and applying the restoration algorithm. For this purpose we will use an image that has been high-pass filtered. The effects of the restoration procedure are easier to evaluate when the image has been enhanced in this way, because it acts primarily on fine detail. The low spatial frequency region often dominates the appearance of imagery, and only by reducing its relative importance to the image can improvements in high-frequency detail be examined closely. Also, local filtering techniques should, in theory, work better on images with small correlation distances. Local techniques ignore information that is far from the area that is restored; to disregard this information is surely justified when distant regions are uncorrelated. Excluding distant regions may be justifiable more generally. We have compared the results of locally restoring the original image with those obtained by restoring the high-frequency portion and recombining it with the low-frequency portion and found no discernible differences. Further investigation of this possible equivalence is necessary to reach reliable conclusions.

*Work performed under the auspices of the US Department of Energy, Contract No. W/405-ENG-50.

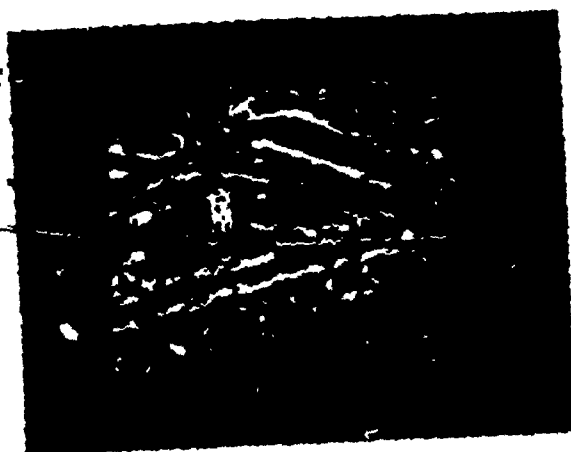


Fig. 1. Image that will be blurred and restored. This has been high-frequency boosted.



Fig. 2. Version of Fig. 1 that was blurred with a 3 by 3 point spread function.

Figure 1 is a 128 by 128 image that has been subjected to a high-frequency boosting, and Fig. 2 is a version of it that was blurred by a point spread function that approximated a circle with a diameter of 3 pixels. We will compute both a global Wiener filter restoration with a local restoration of this blurred image. The Wiener filter is represented by

$$F(f) = \frac{\hat{S}^*(f)\hat{A}(f)}{|\hat{S}(f)|^2\hat{A}(f) + \langle n^2 \rangle} \quad (1)$$

where $\hat{S}(f)$ is the Fourier transform of the point spread function, $\hat{A}(f)$ is an estimate of the Wiener spectrum of the object, and $\langle n^2 \rangle$ is the noise parameter. Equation (1) is the filter that, on the average, minimizes the squares of the differences between the estimated and actual objects. In the usual derivation of Eq. (1), the noise parameter is an assumed noise variance in the image, but better restorations can be computed in practice by treating $\langle n^2 \rangle$ as a free parameter. Usually a value of $\langle n^2 \rangle$ that is smaller than the noise variance results in restorations with sharper edges and enhanced detail, but with more prominent noise patterns. These are often preferred when the noise pattern can be recognized as such. When $\langle n^2 \rangle$ is a free parameter, Eq. (1) has been called the "parametric Wiener filter."

No noise was added to Fig. 2 after it was blurred. If the noise parameter is zero, the filter $F(f)$ has singular points, and we must set $\langle n^2 \rangle$ equal to a small value to make it possible to compute $F(f)$. The result of filtering by this $F(f)$ is Fig. 3. The Fourier transform of $F(f)$, the point spread function of the restoration process, is Fig. 4. It extends 15 to 20 pixels from its center, and pixels this far away from the point being restored influence its estimated value.

Local Filtering

The least squares restoration of an object is a weighted sum of a finite number of nearby image points

$$\hat{o}(\vec{x}_i) = \sum_j c_{ij} i(\vec{x}_j) \quad (2)$$

where the c_{ij} are chosen to minimize the expectation of the square of the difference between the estimated and actual object. The quantity that is minimized is ϵ_i^2 , given by

$$\epsilon_i^2 = \langle |o(\vec{x}_i) - \sum_j c_{ij} i(\vec{x}_j)|^2 \rangle \quad (3)$$



Fig. 3. Conventional Wiener filter restoration of Fig. 2.

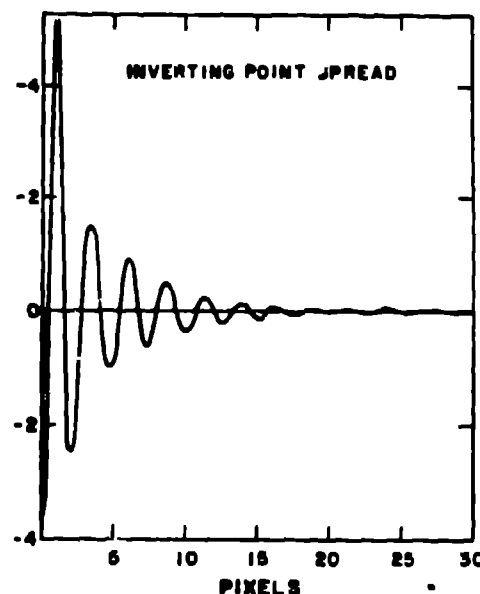


Fig. 4. Point spread function (Fourier transform of filter) used to compute Fig. 3.

The derivative of ϵ_i^2 with respect to each c_{ij} is found and the c_{ij} that minimize ϵ_i^2 are computed by setting this derivative equal to zero. These are functions of the correlation coefficients $\langle o(\vec{x}_i) i(\vec{x}_j) \rangle$ between the object and image at different points and $\langle i(\vec{x}_i) i(\vec{x}_j) \rangle$ between spatially separated values of the image. They are computed by solving the equation

$$\sum_j c_{ij} \langle i(\vec{x}_j) i(\vec{x}_k) \rangle = \langle o(\vec{x}_i) i(\vec{x}_k) \rangle. \quad (4)$$

The correlation coefficients can be expressed in terms of the autocorrelation function $A(\vec{x})$ of the object, the point spread function $s(\vec{x})$ of the imaging system, and the variance $\langle n^2 \rangle$ of the noise. The relations are

$$\langle i(\vec{x}_j) i(\vec{x}_k) \rangle = \int s(\vec{x}' - \vec{x}_j) s(\vec{x}'' - \vec{x}_k) A(\vec{x}' - \vec{x}'') d\vec{x}' d\vec{x}'' + \langle n^2 \rangle \delta_{jk} \quad (5)$$

and

$$\langle o(\vec{x}_i) o(\vec{x}_k) \rangle = \int A(\vec{x}_i - \vec{x}') s(\vec{x}' - \vec{x}_k) d\vec{x}'. \quad (6)$$

We note that the correlation coefficients of widely separated points are small if $A(\vec{x})$ and $s(\vec{x})$ are narrow. We will estimate objects in cases for which they are narrow enough that pixels separated by more than 8 pixels have small correlation coefficients. This should mean that pixels this far away from a point being restored need not be included in the computations. We will solve Eq. (4) for one particular value of the index "i," and use this to estimate $o(\vec{x}_i)$ from the image values $i(\vec{x})$ in a 17 by 17 region that has $o(\vec{x}_i)$ at its center. This region contains 289 pixels.

The solution of Eq. (4) for a region containing 289 image points in general requires the inversion of a 289 by 289 matrix. The size of this matrix can be reduced when the point spread function is symmetric about its horizontal and vertical axes. When this is so, the 17 by 17 pixel region can be divided into zones such that the horizontal and vertical distance of each point in a zone from the center of the region is the same. Figure 14 is a diagram of the first few zones; the region is divided into square pixels that are labeled

such that pixels with the same number comprise a single zone. We estimate the value of the center pixel labeled "i."

As a consequence of the horizontal and vertical symmetry of the point spread function, the correlation coefficients between zones can be computed as a sum of correlation coefficients between individual pixels. For example, the sum of the correlation coefficients between a pixel in zone 5 and all pixels in zone 3 is the same for all pixels in zone 5. This means that members of Eq. (4) can be grouped to form a set of equations that describe the relationships among zones. We use the coefficients to compute the estimated value of zone 1, which consists only of the point that we estimate. There are 45 zones in our 17 by 17 region. The inversion of a 45 by 45 matrix can be performed quickly on many computers.

Figure 5 is a local restoration of Fig. 2 that was performed in the direct domain as described above. It is slightly smaller than Fig. 3, because we made no attempt to restore the 8 points that are closest to the edges. That it is practically indistinguishable from Fig. 3 indicates that this local processing technique is a satisfactory alternative for the global Fourier domain method.

It is of interest to compare the weights c_{ij} that are used in Eq. (1) with the point spread function that was used to produce Fig. 3. Figure 6 is a plot of the average weight (weight per pixel) of each zone as a function of the distance of each zone from zone 1. The plot has many of the same features as the point spread function, i.e., it has lobes at roughly the same distances apart and the same size as the point spread. However, a number of points are displaced from what might be a smooth curve. These differences are important. If the weights of each zone are replaced by the values of the point spread function that are appropriate for their distances from the center, the restoration that results is very poor.

Correlation Function Selection

An estimate of the autocorrelation function $A(\vec{x})$ of the object is used in the computations of the correlation coefficients of Eqs. (5) and (6). The direct-domain equation for the object estimate is similar to Eq. (1) in the sense that this function is multiplied by something that is nearly the same as its inverse when the noise estimate is small. Also, $A(\vec{x})$ appears in Eqs. (5) and (6) under integrals whose values do not depend significantly on small fluctuations in $A(\vec{x})$. For these reasons, good restorations can be computed for a wide range of $A(\vec{x})$.

Nevertheless, values of $A(\vec{x})$ that are unreasonable can adversely affect the quality of the restoration. When significant detail is present in the image, we have used the autocorrelation function estimate



Fig. 5. Local restoration of Fig. 2.

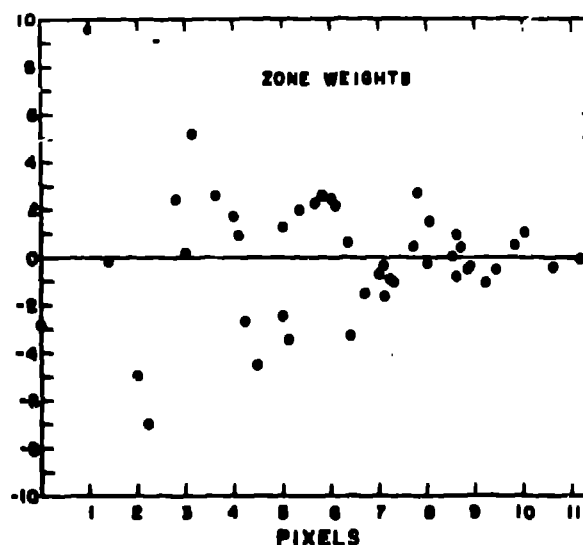


Fig. 6. Average weight of each zone as a function of the distance from zone 1.

$$\begin{aligned}
 A(\vec{x}) &= \exp(-2\pi 10|\vec{x}|/128) & \text{if } |\vec{x}| \leq 6 \\
 &= 0 & \text{if } |\vec{x}| > 6
 \end{aligned}
 \tag{7}$$

where $|\vec{x}|$ is a number of pixels. This exponential autocorrelation function is typical of images that consist of pulses of random height and length, as might occur when many edges are present.² The half-width of this autocorrelation function corresponds to the expected frequency content of the objects whose images have been high-pass filtered as described above. Where there is relatively little detail, a broader autocorrelation is more appropriate. When the autocorrelation function is too broad (or a constant) the matrix that must be inverted to solve Eq. (4) is ill-conditioned, and noise in the image is severely amplified. An autocorrelation function of the form

$$\begin{aligned}
 A(\vec{x}) &= 1 - \frac{1}{4} |\vec{x}| & \text{if } |\vec{x}| \leq 4 \\
 &= 0 & \text{if } |\vec{x}| > 4
 \end{aligned}
 \tag{8}$$

is broader than the exponential function of Eq. (7), yet has not led to the problems associated with ill-conditioning. We have found it useful for images with relatively little detail. When significant detail is present, it does not lead to restorations that are as sharp as those produced using Eq. (7).

Besides the restriction that the assumed autocorrelation function not be broad, there are some other limitations on its choice that we should note. If the function is not circularly symmetric, horizontal and vertical artifacts may appear. This also occurs if the function of Eq. (7) is not set equal to zero for $|\vec{x}| > 6$, even though it is small for larger values of $|\vec{x}|$. Also, if the autocorrelation function is very narrow, e.g., a delta function, the restoration has a granular appearance. Some empirical investigation may be necessary to determine the best $A(\vec{x})$ for a particular type of imagery.

Variable Imaging Parameters

Besides offering computational simplicity, local filtering makes it possible to adjust filter parameters to fit local variations in the characteristics of an image. These variations may arise from a spatially variant point spread function, from signal-dependent noise, or from a nonhomogeneous object. We will demonstrate how the last of these may be treated with an object autocorrelation function and a noise parameter that varies from pixel to pixel. We will restore with a binary filter. We will use the correlation function of Eq. (7) and a noise parameter of 0.006 when the signal-to-noise ratio is above a given threshold, and we will use the autocorrelation function of Eq. (8) and a noise parameter of 0.024 when it is below the threshold. Restoration with a binary filter could be done globally, by computing two restorations and choosing pixels from one or the other according to the same threshold criterion. However, the local technique is not only easier, but it could be extended so that many more than two choices may be made. On a global scale, such an extension is very difficult, if not completely unfeasible.

Figure 7 is an image blurred with a circular point spread function with a diameter of 5 pixels. Gaussian noise is added to the picture that has a variance equal to 0.2 times the variance of the unblurred image. Figure 8 is a restoration which uses the high-signal filter that is described above, and Fig. 9 is a restoration that uses the low-signal filter. Figure 8 shows a granular effect that is typical of Wiener-filter restorations; it occurs because noise is amplified in the restoration process. This effect is not present in Fig. 9, but the restoration is not as sharp.

Figure 10 is the result of a use of a binary filter. The square of each pixel of Fig. 7 was computed and those whose values were less than 0.75 of the variance of the entire image were recomputed with the low-signal parameters. Figure 11 shows in white those pixels that were restored using the high-signal filter.

We note that many of the pixels that do not look like the imagery in their surrounding region were restored as high-signal points while points nearby were restored as low-signal points. Isolated points such as this have been identified as noise with good results.³ Figure 12 was produced from Fig. 10 by locating all points restored with the high-signal parameters with not more than one neighbor that was restored in this way and recomputing these pixels with the low-signal parameters. Figure 13 shows in white the pixels that were processed with high-signal parameters. The image of Fig. 12 is practically as sharp as the restoration of Fig. 8 but lacks its granularity.



Fig. 7. Version of Fig. 1 that was blurred with a 5 by 5 point spread function to which noise was added.



Fig. 8. Figure 7 restored with a filter that is intended for images with high signal-to-noise ratios.



Fig. 9. Figure 7 restored with a filter that is intended for images with low signal-to-noise ratios.



Fig. 10. Figure 7 restored with a combination of the filters used to restore Figs. 8 and 9.

Conclusions

If an estimate of the object autocorrelation function is chosen properly, least squares restoration can be performed in the direct domain with the same results as in the Fourier domain. Such restoration is local in the sense that the restoration of each point depends on the values of a few nearby points and not on distant ones. If the point spread function is circularly symmetric, the autocorrelation function must be too. The autocorrelation function must approach zero smoothly, and it must be a few pixels wide. The local nature of this restoration proves to be advantageous because it permits adjustments to be made to fit spatially variant image conditions such as variant signal-to-noise ratios.

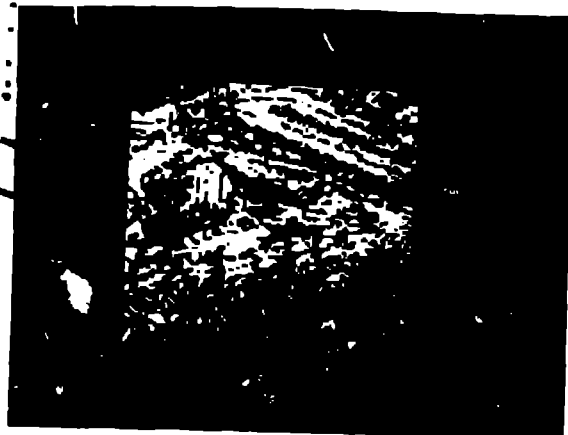


Fig. 11. Those pixels of Fig. 10 that were restored with the high signal-to-noise filter are white, and those restored with the low signal-to-noise filter are black.



Fig. 12. Figure 7 restored with a combination of the high and low signal-to-noise filters. In this restoration, isolated pixels with high values were identified as noise.



Fig. 13. Those pixels of Fig. 12 that were restored with the high signal-to-noise filter are white and the others are black.

10	9	8	7	8	9	10
9	6	5	4	5	6	9
8	5	3	2	3	5	8
7	4	2	1	2	4	7
8	5	3	2	3	5	8
9	6	5	4	5	6	9
10	9	8	7	8	9	10

Fig. 14. Zone diagram. Pixels labeled with the same number comprise a single zone.

References

1. H. C. Andrews and B. R. Hunt, Digital Image Restoration (Prentice-Hall, Englewood Cliffs, NJ, 1977) p 150.
2. S. Goldman, Information Theory (Prentice-Hall, New York, 1953) p 252.
3. H. J. Zweig, E. B. Barrett, and P. C. Hu, "Noise Cheating Image Enhancement," J. Opt. Soc. Am., 65, 1347-1353 (1975).